Scientific Data Visualization Using Hexcore Resources at ORNL

Ross J. Toedte
National Center for Computational Sciences
Oak Ridge National Laboratory

Hexcore Workshop 2010/05/12

Outline

- Visualization centers at ORNL
  - Per-center resources and usage
- Success stories
- Science communication
- Segue to focus on ultra-scale visualization

Visualization Centers at ORNL

- NICS visualization resources
- RDAV
- NCCS Visualization Task Group

NICS Visualization Resources

- Hardware
  - Verne
    - 5-node, 80-core
    - 640GB aggregate memory
- Software
  - plplot, idl, vapor, visit, paraview, vmd, vtk, pgplot
  - lots of ancillary modules for data formats, etc

NICS Visualization Resources

- Usage cases
  - many needs can be met with 1 node
  - significant per-node memory
  - modest parallelism
- Access
  - submit ticket for access to Verne
  - submit ticket to describe your needs

RDAV

- The University of Tennessee Center for Remote Data Analysis and Visualization (RDAV)
- TeraGrid visualization resource center
- participants: ORNL, UT/K, LBNL, UWis, NCSA
- Sean Ahern - PI

Providing analysis services for TeraGrid XD users

- Provide remote and shared resources for the purpose of exploring/analyzing/ visualizing large scale data.
- Provide the ability to easily take advantage of remote and shared computing/ data storage infrastructure.
- Provide unique architecture for data analysis and visualization
- Leverage large amount of existing experience in deploying similar capabilities.
- Allocated through TRAC

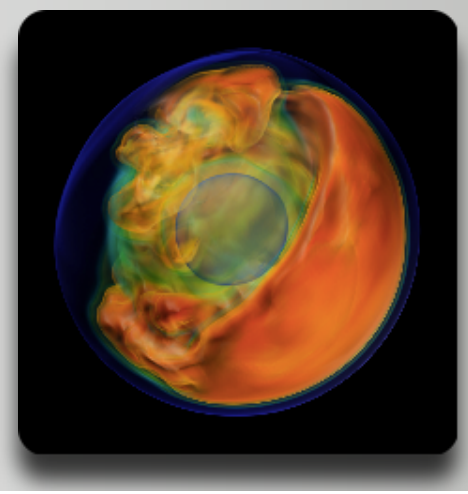

A large SMP is central to RDAV's hardware

- SGI UltraViolet system
  - 1,024 cores (Intel Nehalem EX)
  - 4 TB Global Shared Memory
  - 8-16 NVIDIA Fermi Tesla GPUS – "S" config
- ~1 PB shared filesystem
- ~30 GB/s bandwidth

Early UltraViolet Test Systems

- SGI has delivered a UV test unit to ORNL
- System configuration:
  - 96 Nehalem EX cores
  - 96 GB memory
  - 2x 10 Gigabit Ethernet
  - 2x QDR Infiniband
  - 1x NVIDIA graphics card
- Beginning friendly user access in the next two weeks
- Second test system will be delivered in early May.

- System configuration

es

et

Final UV System – Nautilus

- The final UV system is expected to ship in early June. This will be a forklift upgrade (i.e. the early test system will be physically removed).
- System configuration of Nautilus:
  - 64 blades
    - 128 sockets/1024 cores
    - 4 TB memory
    - SLES 11
  - 4x 10 Gigabit Ethernet
  - 24x QDR 4x Infiniband
  - 8-16 NVIDIA Fermi Tesla
  - ~1 PB parallel file system
- Must be accepted and in production by early July

How Nautilus Fits Into NICS

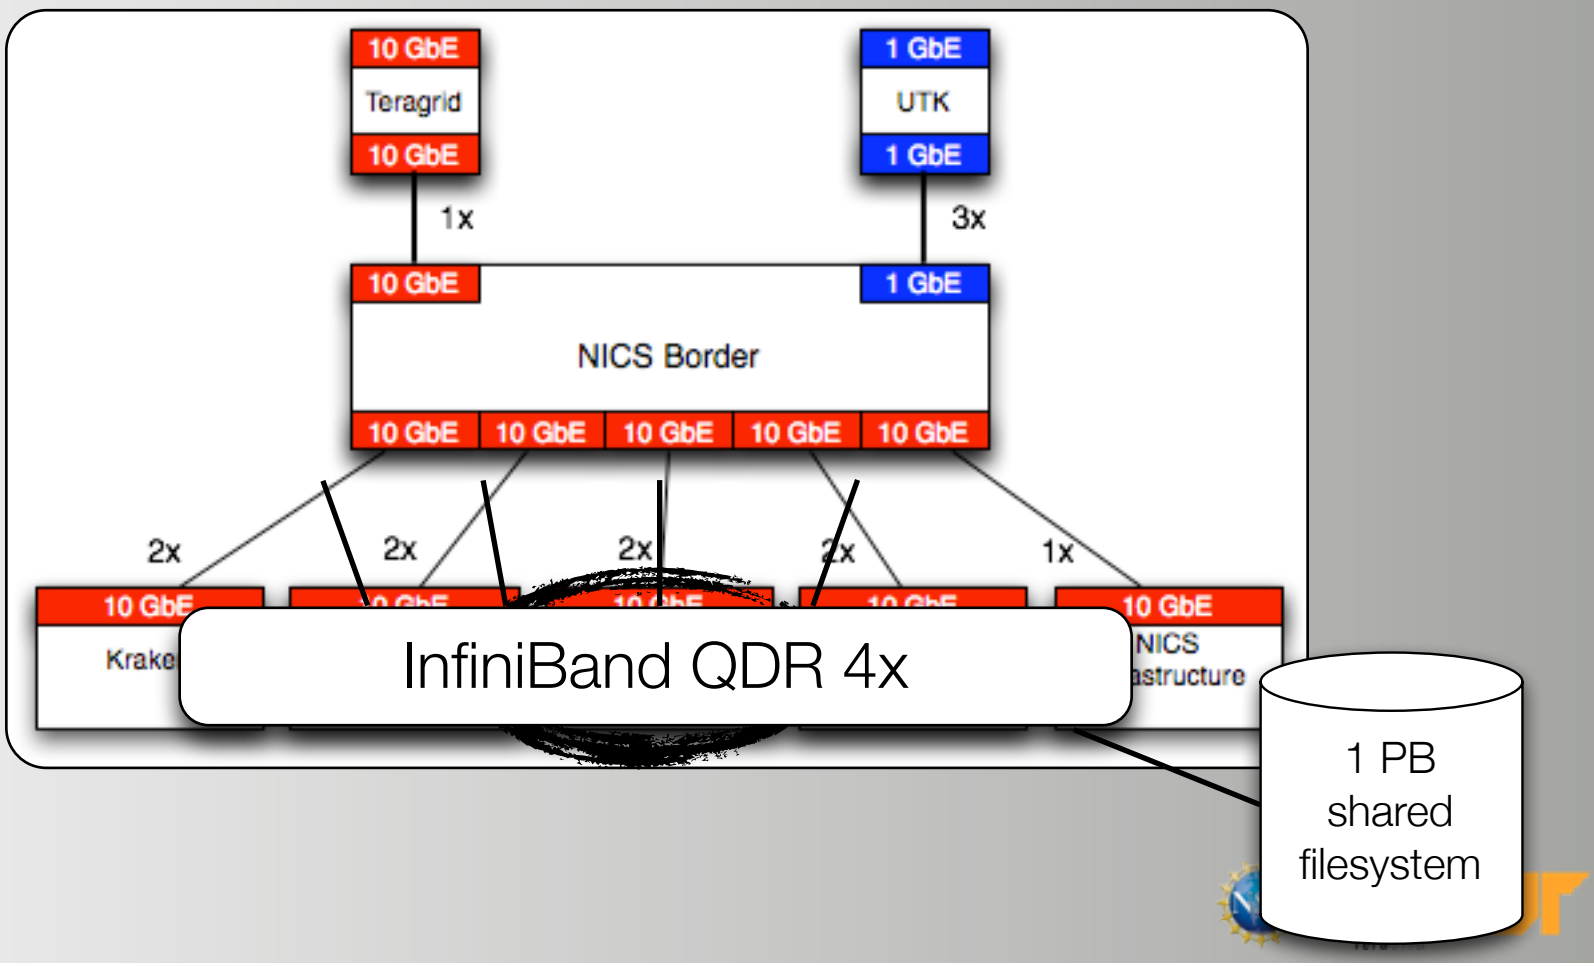

RDAV provides a range of software services

- Analysis applications: to be dictated by user needs and technology needed to solve user problems. “Whatever it takes!”
- Remote visualization and image generation
  - Provide interactive and batch image generation tools. (gnuplot, ImageMagick, etc.)
  - Remote parallel visualization (VisIt, ParaView, etc.)
  - Tools for custom application development
- Data analysis and statistical analysis
  - Octave, Parallel R, Matlab, etc.
- Workflow systems
  - DAGMan system automates batch actions on behalf of users
  - Infrequent current use, however, value is increasing and many users wish to explore.
- Dashboard delivery
  - Leverage DoE funding for eSimMon dashboard system.
- Portal system
  - Builds upon standard Liferay platform
  - Provides SAS services for analysis and visualization

NCCS Visualization Task Group

- Principle focus on DOE/INCITE science teams and discretionary allocations
- diverse applications science experience
- pragmatic approach to working with science teams
  - match skills, experience to project needs
  - project team articulates needs/lack thereof
  - "It's your piece of layer cake."

NCCS Visualization Task Group

- Personnel
  - Sean Ahern (task lead)
  - George Ostrouchov (statistics)
  - David Banks (UT/K, novel data/representations)
  - Jamison Daniel (climate, GPU)
  - Mike Matheson (combustion, Blender)
  - Dave Pugmire (fusion)
  - Jeremy Meredith (0.5, GPU, materials)
  - Rob Sisneros (post-doc, summary vis)
  - Ross Toedte (astro)

NCCS Visualization Task Group

- Personnel
  - Sean Ahern - ahern@ornl.gov
  - George Ostrouchov - ost@ornl.gov
  - David Banks - banksdc@ornl.gov
  - Jamison Daniel - d65@ornl.gov
  - Mike Matheson - 5iv@ornl.gov
  - Dave Pugmire - dpn@ornl.gov
  - Jeremy Meredith - js9@ornl.gov
  - Rob Sisneros - vso@ornl.gov
  - Ross Toedte - rjt@ornl.gov

NCCS Visualization Task Group

- Hardware
  - Lens
    - 32-node, 512-core
    - 2TB aggregate memory
    - 1 NVidia Tesla C1060 and 1 GTX8800 / node
  - EVEREST cluster
    - 16-node, 32-core
    - dedicated to driving EVEREST PowerWall
  - EVEREST PowerWall

NCCS Visualization Task Group

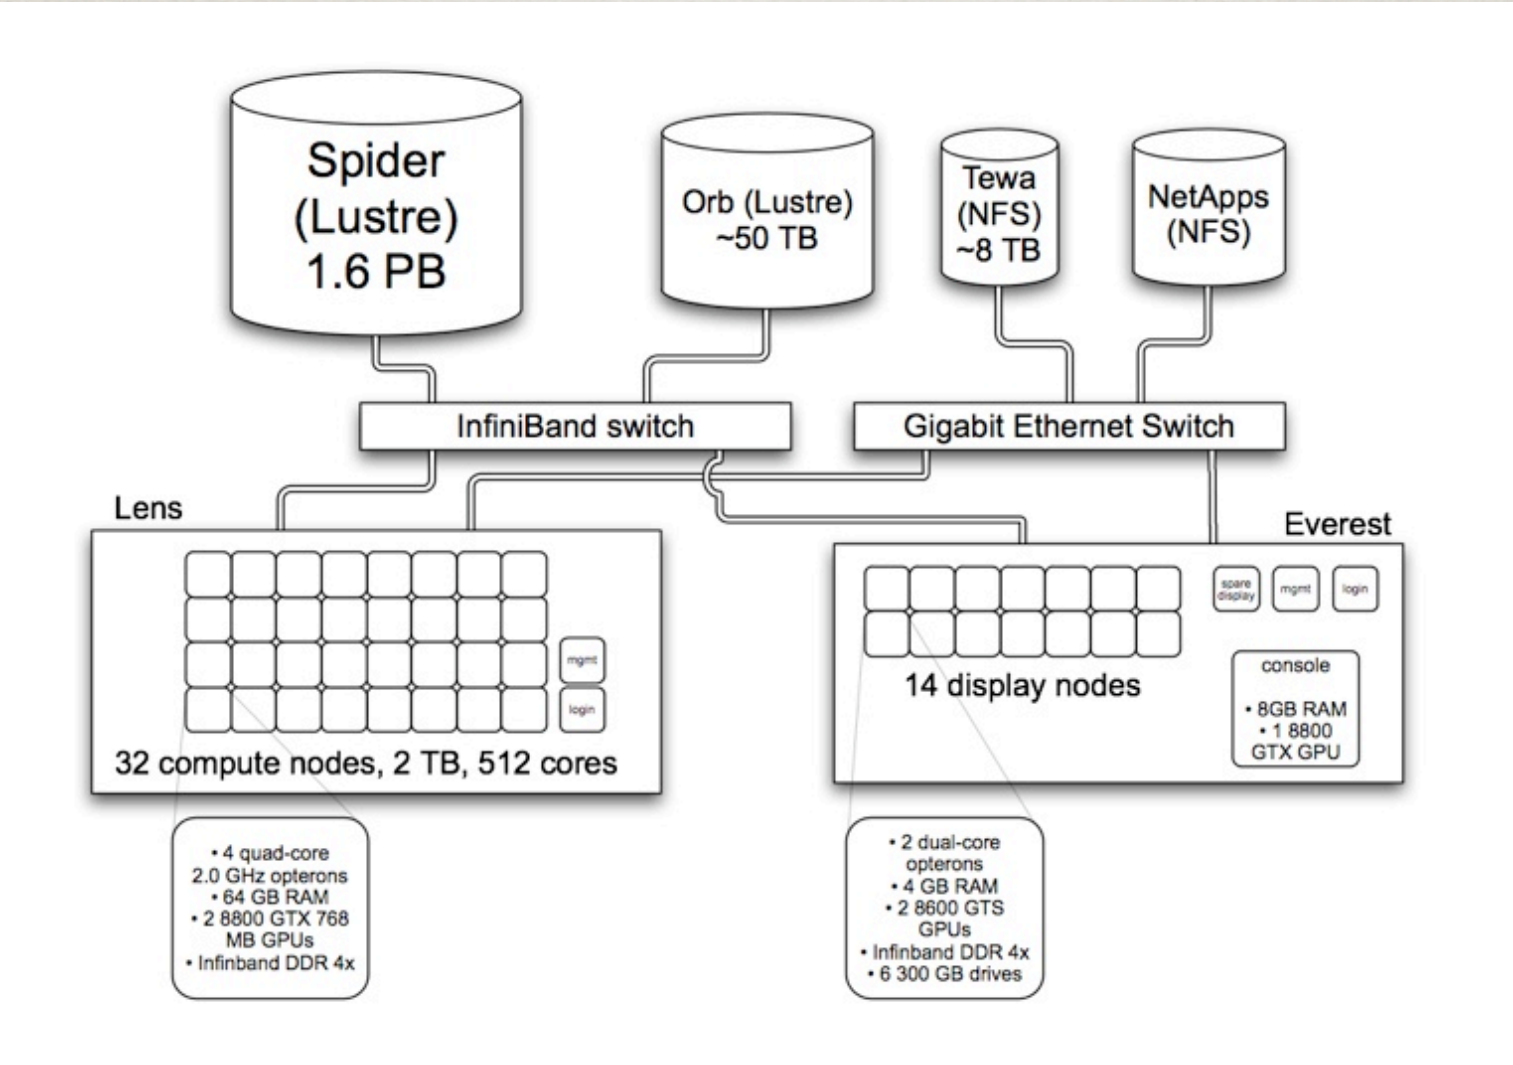

NCCS Visualization Task Group

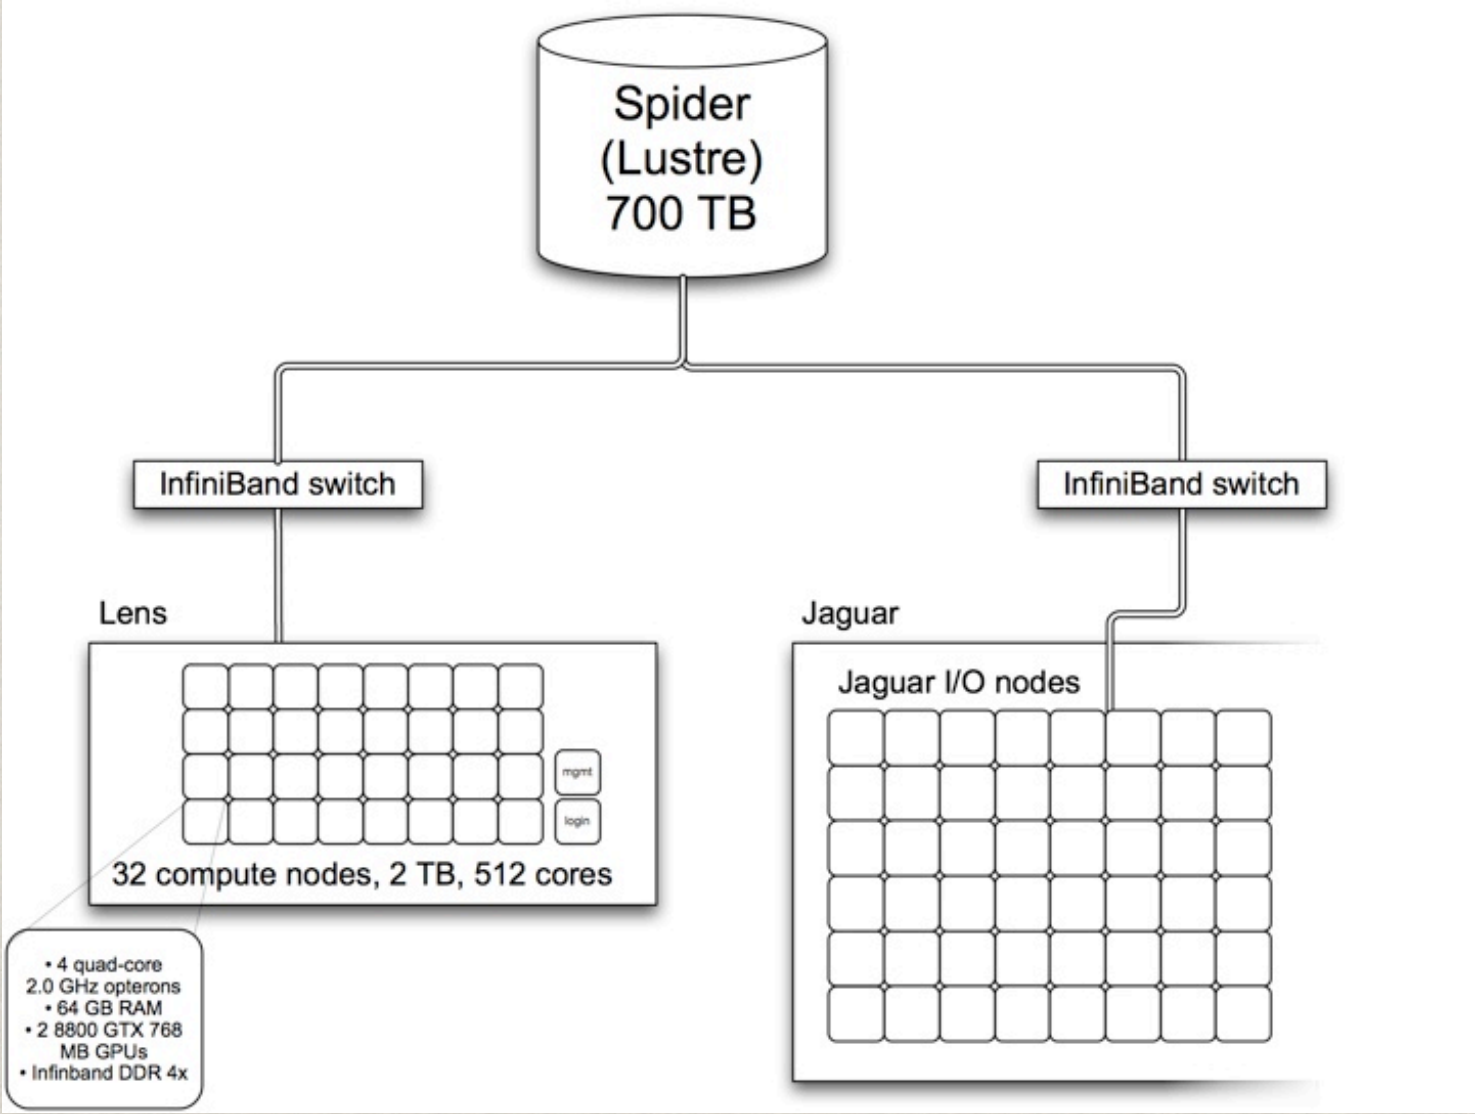

NCCS Visualization Task Group

- visualization software
  - avs-express
  - ensight
  - ferret
  - gnuplot
  - grads
  - idl
  - ncview
  - pgplot
  - povray
  - tecplot
  - visit
  - vmd
- related software
  - adios
  - chromium
  - cuda
  - everest
  - ggobi
  - hdf5
  - java-jdk
  - java-jre
  - matlab
  - ncl
  - nco
  - netcdf
  - p-netcdf
  - r
  - silo

NCCS Visualization Task Group

- Scientific Linux OS
- Queueing using PBS+Moab; soon to be SLURM
- Modules for viz tools and ancillary data needs
- Preferential usage by analysis jobs; make an explicit account request for this

NCCS Visualization Task Group

- EVEREST Large format PowerWall
- Meant for high-resolution and -dimension studies
- Video switcher coming soon

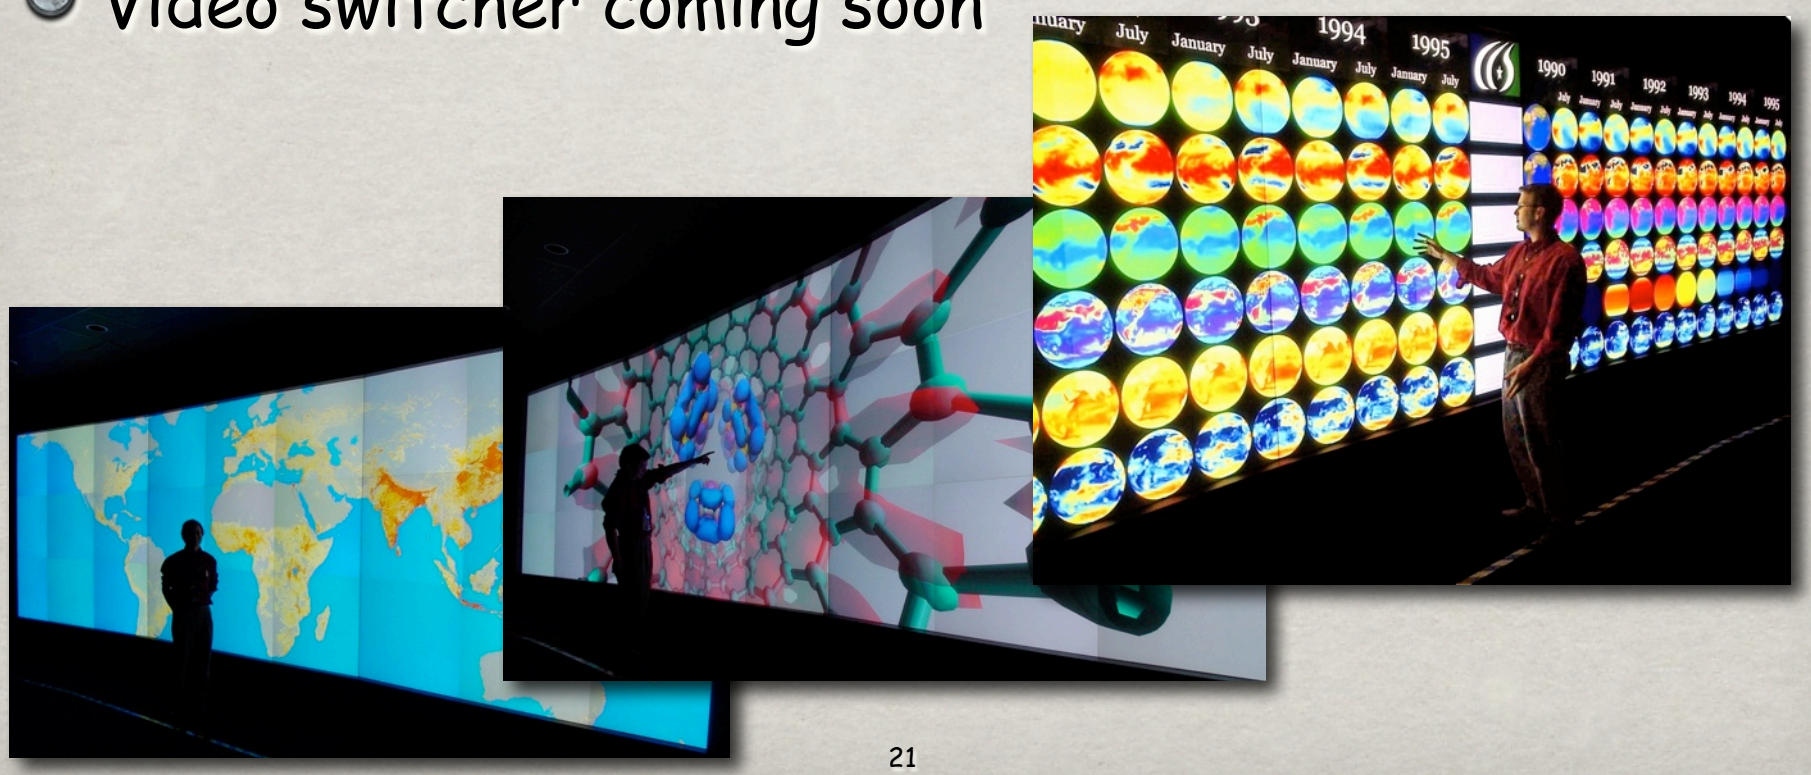

Success stories

- diverse modes of interaction between applications teams and visualization team
- "Whatever it takes" attitude, but requires communication between application team and visualization team
  - considers existing communication fabric
  - considers domain-specific morés
  - usually involves science orientation for viz staff

HOMME-TMQ (Water Vapor)

- Daniel (ORNL), Taylor (SNL), Evans (ORNL), Hack (ORNL), 2009
- embedded visualization staff

Coal Gasifier

- Syamlal (NETL), Gel (NETL, Alpemi), Toedte (ORNL), 2010
- template and workflow development for use by science team
- final compositing by ORNL staff

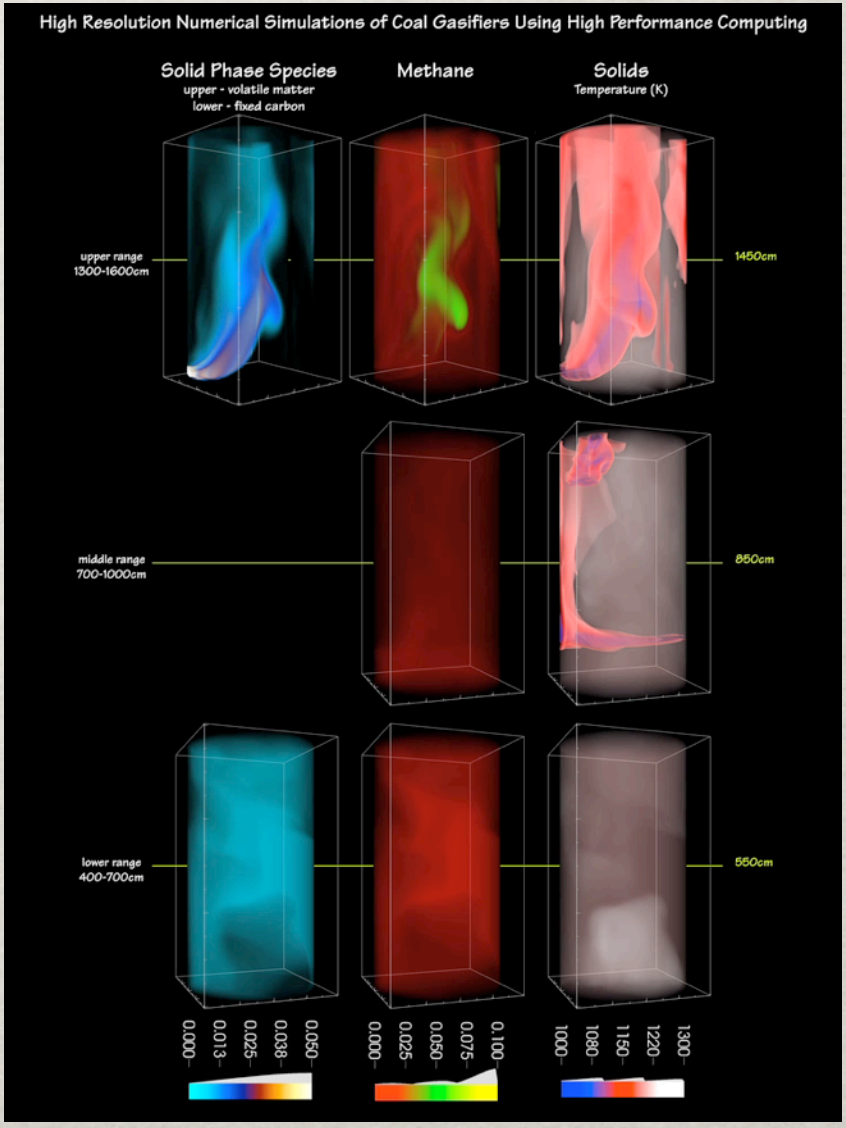

Via Lactae

- Madau, Diemand, et al (UCO-Lick)
- ORNL viz team prototyped VisIt exploration
- in the end, science team used own
- shared final anim

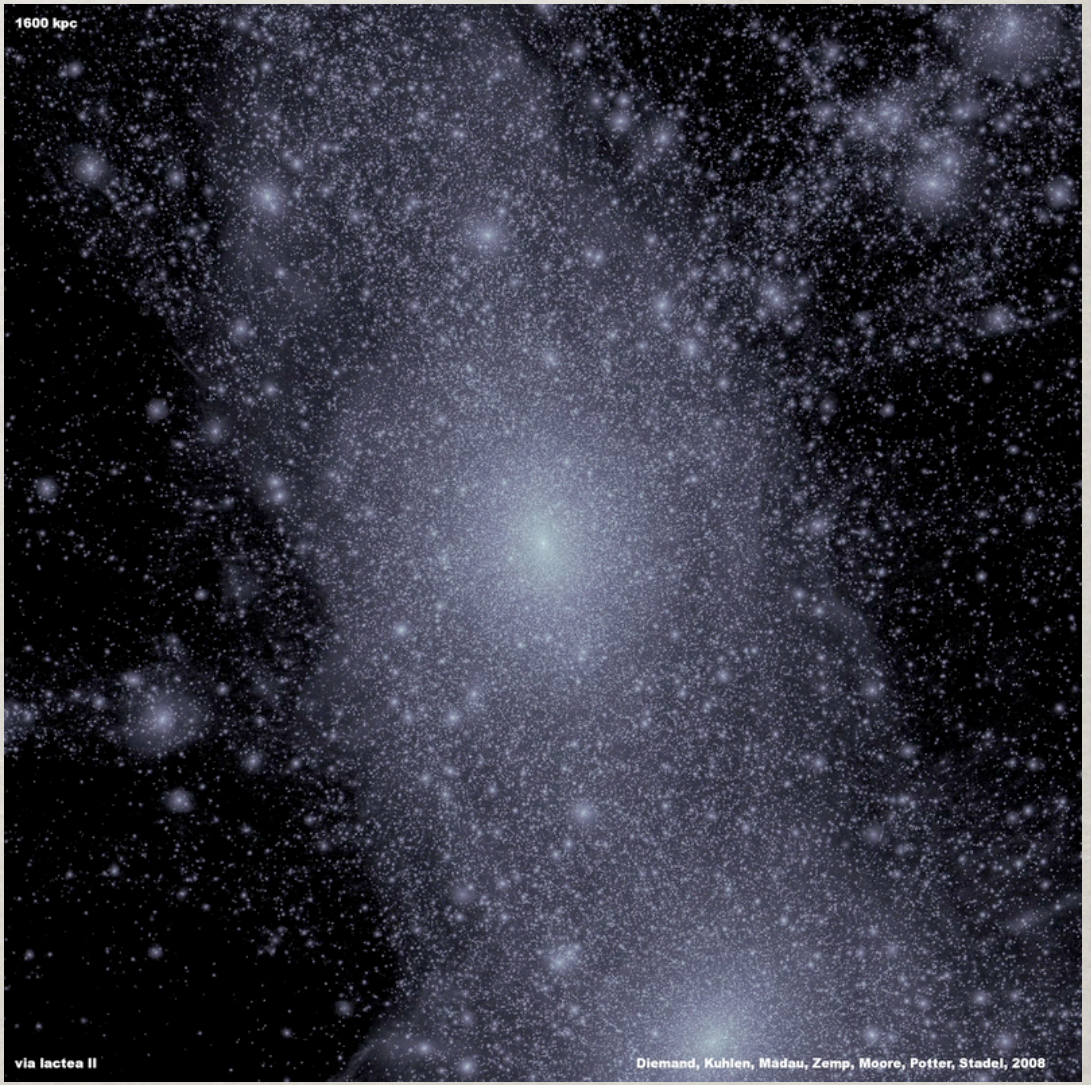

Science Communication

- We give 400-500 tours a year to external visitors
- Visitors include people involved at the highest levels of government, academic and private research
- The "So what?" question needs answering
- Help us present and articulate your work

Ultra-scale visualization

...Dave Pugmire